

Languages cool as they expand: Allometric scaling and the decreasing need for new words

Alexander M. Petersen,^{1,*} Joel N. Tenenbaum,^{2,3} Shlomo Havlin,⁴ H. Eugene Stanley,² and Matjaž Perc^{5,†}

¹Laboratory for the Analysis of Complex Economic Systems,
IMT Lucca Institute for Advanced Studies, Lucca 55100, Italy

²Center for Polymer Studies and Department of Physics,
Boston University, Boston, Massachusetts 02215, USA

³Operations and Technology Management, School of Management,
Boston University, Boston, Massachusetts 02215, USA

⁴Minerva Center and Department of Physics, Bar-Ilan University, Ramat-Gan 52900, Israel

⁵Department of Physics, Faculty of Natural Sciences and Mathematics,
University of Maribor, Koroška cesta 160, SI-2000 Maribor, Slovenia

We analyze the occurrence frequencies of over 15 million words recorded in millions of books published during the past two centuries in seven different languages. For all languages and chronological subsets of the data we confirm that two scaling regimes characterize the word frequency distributions, with only the more common words obeying the classic Zipf law. Using corpora of unprecedented size, we test the allometric scaling relation between the corpus size and the vocabulary size of growing languages to demonstrate a decreasing marginal need for new words, a feature that is likely related to the underlying correlations between words. We calculate the annual growth fluctuations of word use which has a decreasing trend as the corpus size increases, indicating a slowdown in linguistic evolution following language expansion. This “cooling pattern” forms the basis of a third statistical regularity, which unlike the Zipf and the Heaps law, is dynamical in nature.

Books in libraries and attics around the world constitute an immense “crowd-sourced” historical record that traces the evolution of culture back beyond the limits of oral history. However, the disaggregation of written language into individual books makes the longitudinal analysis of language a difficult open problem. To this end, the book digitization project at Google Inc. presents a monumental step forward providing an enormous, publicly accessible, collection of written language in the form of the *Google Books Ngram Viewer* web application [1]. Approximately 4% of all books ever published have been scanned, making available over 10^7 occurrence time series (word-use trajectories) that archive cultural dynamics in seven different languages over a period of more than two centuries. This dataset highlights the utility of open “Big Data,” which is the gateway to “metaknowledge” [2], the knowledge about knowledge. A digital data deluge is sustaining extensive interdisciplinary research efforts towards quantitative insights into the social and natural sciences [3–7].

“Culturomics,” the use of high-throughput data for the purpose of studying human culture, is a promising new empirical platform for gaining insight into subjects ranging from political history to epidemiology [8]. As first demonstrated by Michel et al. [8], the *Google* n-gram dataset is well-suited for examining the microscopic properties of an entire language ecosystem. Using this dataset to analyze the growth patterns of individual word frequencies, Petersen et al. [9] recently identified tipping points in the life trajectory of new words, statistical patterns that govern the fluctuations in word use, and quantitative measures for cultural memory. The statistical properties of cultural memory, derived from the quantitative analysis of individual word-use trajectories, were also inves-

tigated by Gao et al. [10], who found that words describing social phenomena tend to have different long-range correlations than words describing natural phenomena.

Here we study the growth and evolution of written language by analyzing the macroscopic scaling patterns that characterize word-use. Using the *Google* 1-gram data collected at the 1-year time resolution over the period 1800–2008, we quantify the annual fluctuation scale of words within a given corpora and show that languages can be said to “cool by expansion.” This effect constitutes a dynamic law, in contrast to the static laws of Zipf and Heaps which are founded upon snapshots of single texts. The Zipf law [11–17], quantifying the distribution of word frequencies, and the Heaps law [13, 18–20], relating the size of a corpus to the vocabulary size of that corpus, are classic paradigms that capture many complexities of language in remarkably simple statistical patterns. While these laws have been exhaustively tested on relatively small snapshots of empirical data, here we test the validity of these laws using extremely large corpora.

Interestingly, we observe two scaling regimes in the probability density functions of word usage, with the Zipf law holding only for the set of more frequently used words, referred to as the “kernel lexicon” by Ferrer i Cancho et al. [14]. The word frequency distribution for the rarely used words constituting the “unlimited lexicon” [14] obeys a distinct scaling law, suggesting that rare words belong to a distinct class. This “unlimited lexicon” is populated by highly technical words, new words, numbers, spelling variants of kernel words, and optical character recognition (OCR) errors.

Many new words start in relative obscurity, and their eventual importance can be under-appreciated by their initial frequency. This fact is closely related to the information cost of introducing new words and concepts. For single topical texts, Heaps observed that the vocabulary size exhibits sub-linear growth with document size [18]. Extending this concept to entire corpora, we find a scaling relation that indicates a decreas-

*Electronic address: alexander.petersen@imtlucca.it

†Electronic address: matjaz.perc@uni-mb.si

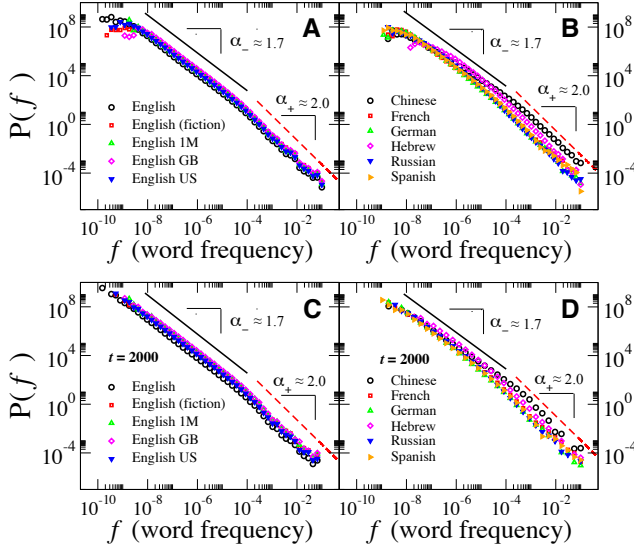


FIG. 1: Two-regime scaling distribution of word frequency. The kink in the probability density functions $P(f)$ occurs around $f_x \approx 10^{-5}$ for each corpora analyzed (see legend). (A,B) Data from all years are aggregated into a single distribution. (C,D) $P(f)$ comprising data from only year $t = 2000$ providing evidence that the distribution is stable even over shorter time frames and likely emerges in corpora that are sufficiently large to be comprehensive of the language studied. For details concerning the scaling exponents we refer to Table I and the main text.

ing “marginal need” for new words which are the manifestation of cultural evolution and the seeds for language growth. We introduce a pruning method to study the role of infrequent words on the allometric scaling properties of language. By studying progressively smaller sets of the kernel lexicon we can better understand the marginal utility of the core words. The pattern that arises for all languages analyzed provides insight into the intrinsic dependency structure between words.

The correlations in word use can also be author and topic dependent. Bernhardsson et al. recently introduced the “metabook” concept [19, 20], according to which word-frequency structures are author-specific: the word-frequency characteristics of a random excerpt from a compilation of everything that a specific author could ever conceivably write (his/her “metabook”) should accurately match those of the author’s actual writings. It is not immediately obvious whether a compilation of all the metabooks of all authors would still conform to the Zipf law and the Heaps law. The immense size and time span of the *Google* n-gram dataset allows us to examine this question in detail.

Results

Longitudinal analysis of written language. Allometric scaling analysis [21] is used to quantify the role of system size on general phenomena characterizing a system, and has been applied to systems as diverse as the metabolic rate of mitochon-

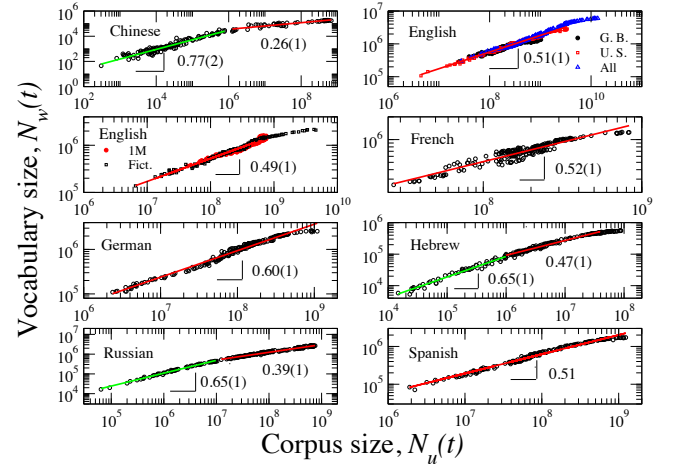


FIG. 2: Allometric scaling of language. Scatter plots of the output corpora size N_u given the empirical vocabulary size N_w using all data ($U_c = 0$) over the 209-year period 1800–2008. Shown are OLS estimation of the exponent b quantifying the Heaps’ law relation $N_w \sim [N_u]^b$.

dria [22] and city growth [23–29]. Indeed, city growth shares two common features with the growth of written text: (i) the Zipf law is able to describe the distribution of city sizes regardless of country or the time period of the data [26], and (ii) city growth has inherent constraints due to geography, changing labor markets and their effects on opportunities for innovation and wealth creation [27, 28], just as vocabulary growth is constrained by human brain capacity and the varying utilities of new words across users [14].

We construct a word counting framework by first defining the quantity $u_i(t)$ as the number of times word i is used in year t . Since the number of books and the number of distinct words grow dramatically over time, we define the *relative* word use, $f_i(t)$, as the fraction of the total body of text occupied by word i in the same year

$$f_i(t) \equiv u_i(t)/N_u(t), \quad (1)$$

where the quantity $N_u(t) \equiv \sum_{i=1}^{N_w(t)} u_i(t)$ is the total number of indistinct word uses while $N_w(t)$ is the total number of distinct words digitized from books printed in year t . Both the N_w (“types” giving the vocabulary size) and the N_u (“tokens” giving the size of the body of text) are generally increasing over time.

The Zipf law and the two scaling regimes. Zipf investigated a number of bodies of literature and observed that the frequency of any given word is roughly inversely proportional to its rank [11], with the frequency of the z -ranked word given by the relation

$$f(z) \sim z^{-\zeta}, \quad (2)$$

with a scaling exponent $\zeta \approx 1$. This empirical law has been confirmed for a broad range of data, ranging from income rankings, city populations, and the varying sizes of

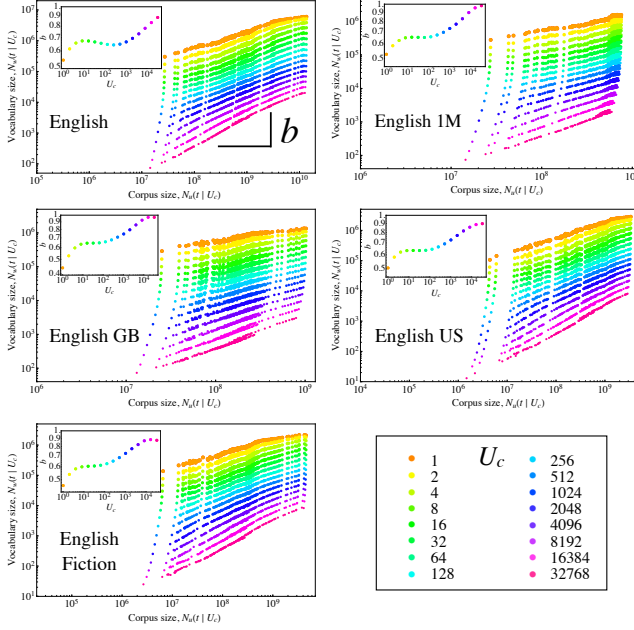


FIG. 3: Pruning reveals the variable marginal return of words. The Heaps scaling exponent b depends on the extent of the inclusion of the rarest words. For a given corpora and U_c value we make a scatter plot between $N_w(t|U_c)$ and $N_u(t|U_c)$ using words with $u_i(t) \geq U_c$. (Panel Inset) We use OLS estimation to estimate the scaling exponent $b(U_c)$ for the model $N_w(t|U_c) \sim [N_u(t|U_c)]^b$ to show that $b(U_c)$ increases from approximately 0.5 towards unity as we prune the corpora of extremely rare words. Our longitudinal language analysis provides insight into the structural importance of the most frequent words which are used more times per appearance and which play a crucial role in the usage of new and rare words.

avalanches, forest fires [30] and firm size [31] to the linguistic features of noncoding DNA [32]. The Zipf law can be derived through the “principle of least effort,” which minimizes the communication noise between speakers (writers) and listeners (readers) [16]. The Zipf law has been found to hold for a large dataset of English text [14], but there are interesting deviations observed in the lexicon of individuals diagnosed with schizophrenia [15]. Here, we also find statistical regularity in the distribution of relative word use for 11 different datasets, each comprising more than half a million distinct words taken from millions of books [8].

Figure 1 shows the probability density functions $P(f)$ resulting from data aggregated over all the years (A,B) as well as over 1-year periods as demonstrated for the year $t = 2000$ (C,D). Regardless of the language and the considered time span, the probability density functions are characterized by a striking two-regime scaling, which was first noted by Ferrer i Cancho and Solé [14], and can be quantified as

$$P(f) \sim \begin{cases} f^{-\alpha_-}, & \text{if } f < f_\times \text{ [“unlimited lexicon”]} \\ f^{-\alpha_+}, & \text{if } f > f_\times \text{ [“kernel lexicon”]}. \end{cases} \quad (3)$$

These two regimes, designated “kernel lexicon” and “unlimited lexicon,” are thought to reflect the cognitive constraints of the brain’s finite vocabulary [14]. The specialized words

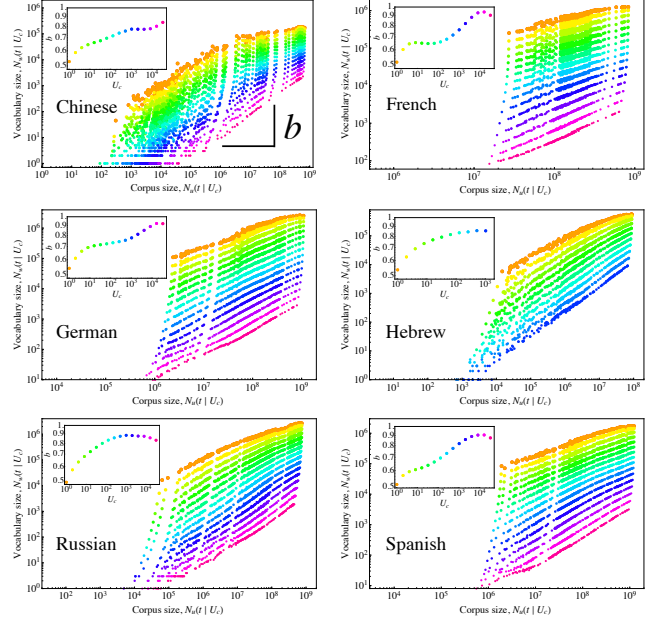


FIG. 4: Pruning reveals the variable marginal return of words. The Heaps scaling exponent b depends on the extent of the inclusion of the rarest words. For a given corpora and U_c value we make a scatter plot between $N_w(t|U_c)$ and $N_u(t|U_c)$ using words with $u_i(t) \geq U_c$, using the same data color- U_c correspondence as in Fig. 3. (Panel Inset) We use OLS estimation to estimate the scaling exponent $b(U_c)$ for the model $N_w(t|U_c) \sim [N_u(t|U_c)]^b$ to show that $b(U_c)$ increases from approximately 0.5 towards unity as we prune the corpora of extremely rare words. Our longitudinal language analysis provides insight into the structural importance of the most frequent words which are used more times per appearance and which play a crucial role in the usage of new and rare words.

found in the unlimited lexicon are not universally shared and are used significantly less frequently than the words in the kernel lexicon. This is reflected in the kink in the probability density functions and gives rise to the anomalous two-scaling distribution shown in Fig. 1.

The exponent α_+ and the corresponding rank-frequency scaling exponent ζ in Eq. (2) are related asymptotically by [14]

$$\alpha_+ \approx 1 + 1/\zeta, \quad (4)$$

with no analogous relationship for the unlimited lexicon values α_- and ζ_- . Table I lists the average α_+ and α_- values calculated by aggregating α_\pm values for each year using a maximum likelihood estimator for the power-law distribution [33]. We characterize the two scaling regimes using a crossover region around $f_\times \approx 10^{-5}$ to distinguish between α_- and α_+ : (i) $10^{-8} \leq f \leq 10^{-6}$ corresponds to α_- and (ii) $10^{-4} \leq f \leq 10^{-1}$ corresponds to α_+ . For the words that satisfy $f \gtrsim f_\times$ that comprise the kernel lexicon, we verify the Zipf scaling law $\zeta \approx 1$ (corresponding to $\alpha \approx 2$) for all corpora analyzed. For the unlimited lexicon regime $f \lesssim f_\times$, however, the Zipf law is not obeyed, as we find $\alpha_- \approx 1.7$. Note that α_- is significantly smaller in the Hebrew, Chinese, and the Russian corpora, which suggests

TABLE I: Summary of the scaling exponents characterizing the Zipf law and the Heaps law. To calculate $\sigma_r(t|f_c)$ (see Figs. 6 and 7) we use only the relatively common words that meet the criterion that their average word use $\langle f_i \rangle$ over the entire word history is larger than a threshold $f_c = 10/\text{Min}[N_u(t)]$ listed in the first column for each corpus. The b values shown are calculated using all words ($U_c = 0$). The “unlimited lexicon” scaling exponent $\alpha_-(t)$ is calculated for $10^{-8} < f < 10^{-6}$ and the “kernel lexicon” exponent $\alpha_+(t)$ is calculated for $10^{-4} < f < 10^{-1}$ using the maximum likelihood estimator method for each year. The average and standard deviation ($\langle \dots \rangle \pm \sigma$) listed are computed using the $\alpha_+(t)$ and $\alpha_-(t)$ values over the 209-year period 1800–2008 (except for Chinese, which is calculated from 1950–2008 data). We show the Zipf scaling exponent calculated as $\zeta = 1/(\langle \alpha_+ \rangle - 1)$. The last column indicates the β scaling exponents from Fig. 7(A).

Corpus (1-grams)	Scaling parameters					
	$\text{Min}[N_u(t)]$	$b(U_c = 0)$	$\langle \alpha_- \rangle$	$\langle \alpha_+ \rangle$	ζ	β
Chinese	35, 394	0.77 ± 0.02	1.49 ± 0.15	1.91 ± 0.04	1.10 ± 0.05	0.20 ± 0.01
English	42, 786, 702	0.54 ± 0.01	1.73 ± 0.05	2.04 ± 0.06	0.96 ± 0.06	0.19 ± 0.01
English fiction	13, 184, 111	0.49 ± 0.01	1.68 ± 0.10	1.97 ± 0.04	1.03 ± 0.04	0.18 ± 0.01
English GB	38, 956, 621	0.44 ± 0.01	1.71 ± 0.07	2.02 ± 0.05	0.98 ± 0.05	0.17 ± 0.01
English US	5, 821, 340	0.51 ± 0.01	1.70 ± 0.08	2.03 ± 0.06	0.97 ± 0.06	0.18 ± 0.01
English 1M	42, 778, 968	0.53 ± 0.01	1.71 ± 0.04	2.04 ± 0.06	0.96 ± 0.06	0.25 ± 0.01
French	34, 198, 362	0.52 ± 0.01	1.69 ± 0.06	1.98 ± 0.04	1.02 ± 0.04	0.26 ± 0.01
German	2, 274, 842	0.60 ± 0.01	1.63 ± 0.16	2.02 ± 0.03	0.98 ± 0.03	0.27 ± 0.01
Hebrew	9, 482	0.47 ± 0.01	1.34 ± 0.09	2.06 ± 0.05	0.94 ± 0.05	0.35 ± 0.01
Russian	6, 944, 366	0.65 ± 0.01	1.55 ± 0.17	2.04 ± 0.06	0.96 ± 0.06	0.08 ± 0.01
Spanish	1, 777, 563	0.51 ± 0.01	1.61 ± 0.15	2.07 ± 0.04	0.93 ± 0.04	0.26 ± 0.01

that a more generalized version of the Zipf law [14] may be needed, one which is slightly language-dependent, especially when taking into account the usage of specialized words from the unlimited lexicon.

The Heaps law and the increasing marginal returns of new words. Heaps observed that vocabulary size, i.e. the number of distinct words, exhibits a sub-linear growth with document size [18]. This observation has important implications for the “return on investment” of a new word as it is established and becomes disseminated throughout the literature of a given language. As a proxy for this return, Heaps studied how often new words are invoked in lieu of preexisting competitors and examined the linguistic value of new words and ideas by analyzing the relation between the total number of words printed in a body of text N_u , and the number of these which are distinct N_w , i.e. the vocabulary size [18]. The marginal returns of new words, $\partial N_u / \partial N_w$ quantifies the impact of the addition of a single word to the vocabulary of a corpus on the aggregate output (corpus size).

For individual books, the empirically-observed scaling relation between N_u and N_w obeys

$$N_w \sim (N_u)^b, \quad (5)$$

with $b < 1$, with Eq. (5) referred to as “the Heaps law”. It has subsequently been found that Heaps’ law emerges naturally in systems that can be described as sampling from an underlying Zipf distribution. In an information theoretic formulation of the the abstract concept of word cost, B. Mandelbrot predicted the relation $b = 1/\zeta$ in 1961 [34], where ζ is the scaling

exponent corresponding to α_+ , as in Eqs. (3) and (4). This prediction is limited to relatively small texts where the unlimited lexicon, which manifests in the α_- regime, does not play a significant role. A mathematical extension of this result for general underlying rank-distributions is also provided by Karlin [35] using an infinite urn scheme, and extended to broader classes of heavy-tailed distributions recently by Gnedin et al. [36]. Recent research efforts using stochastic master equation techniques to model the growth of a book have also predicted this intrinsic relation between Zipf’s law and Heaps’ law [13, 37, 38].

Figure 2 confirms a sub-linear scaling ($b < 1$) between N_u and N_w for each corpora analyzed. These results show how the marginal returns of new words are given by

$$\frac{\partial N_u}{\partial N_w} \sim (N_w)^{(1-b)/b}, \quad (6)$$

which is an increasing function of N_w for $b < 1$. Thus, the relative increase in the induced volume of written languages is larger for new words than for old words. This is likely due to the fact that new words are typically technical in nature, requiring additional explanations that put the word into context with pre-existing words. Specifically, a new word requires the additional use of preexisting words as a result of both (i) the explanation of the content of the new word using existing technical terms, and (ii) the grammatical infrastructure necessary for that explanation. Hence, there are large spillovers in the size of the written corpus that follow from the intricate dependency structure of language stemming from the various grammatical roles [39, 40].

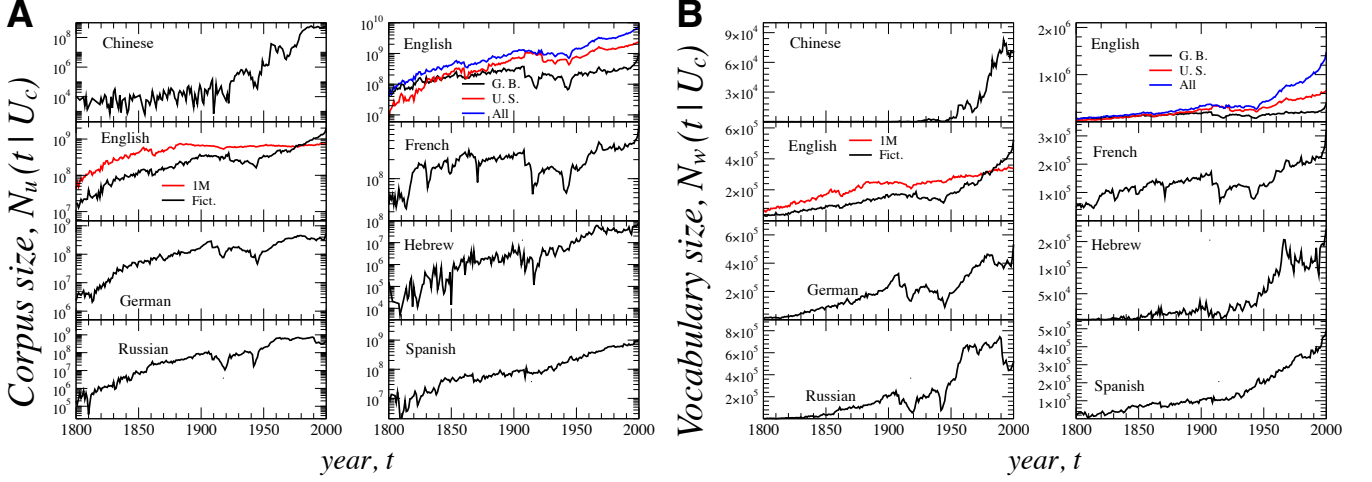


FIG. 5: Literary productivity and vocabulary size in the *Google Inc.* 1-gram dataset over the past two centuries. (A) Total size of the different corpora $N_u(t|U_c)$ over time, calculated by using words that satisfy $u_i(t) \geq U_c \equiv 16$ to eliminate extremely rare 1-grams. (B) Size of the written vocabulary $N_w(t|U_c)$ over time, calculated under the same conditions as (A).

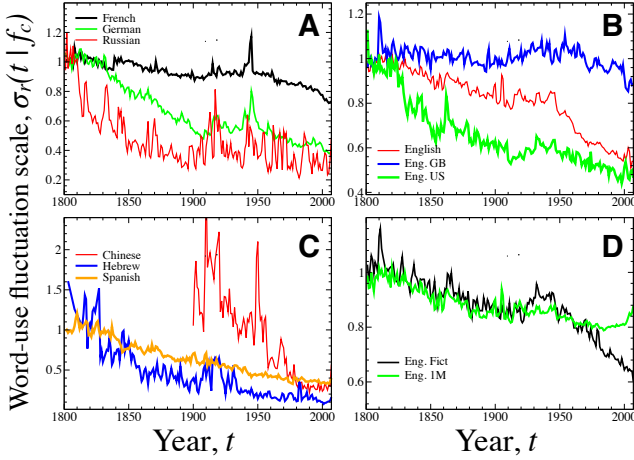


FIG. 6: Non-stationarity in the characteristic growth fluctuation of word use. The standard deviation $\sigma_r(t|f_c)$ of the logarithmic growth rate $r_i(t)$ is presented for all examined corpora. There is an overall decreasing trend arising from the increasing size of the corpora, as depicted in Fig. 5(A). On the other hand, the steady production of new words, as depicted in Fig. 5(B) counteracts this effect. We calculate $\sigma_r(t|f_c)$ using the relatively common words that meet the criterion that their average word use $\langle f_i \rangle$ over the entire word history T_i (using words with lifetime $T_i \geq 10$ years) is larger than a threshold $f_c \equiv 1/\text{Min}[N_u(t)]$ (see Table I).

In order to investigate the role of rare and new words, we calculate N_u and N_w using only words that have appeared at least U_c times. We select the absolute number of uses as a word use threshold because a word in a given year can not appear with a frequency less than $1/N_u$, hence any criteria using relative frequency would necessarily introduce a bias for small corpora samples. This choice also eliminates words that can spuriously arise from Optical Character Recognition (OCR) errors in the digitization process and also from intrinsic

spelling errors and orthographic spelling variations.

Figures 3 and 4 show the relational dependence of N_u and N_w on the exclusion of low-frequency words using a variable cutoff $U_c = 2^n$ with $n = 0 \dots 11$. As U_c increases the Heaps scaling exponent increases from $b \approx 0.5$, approaching $b \approx 1$, indicating that core words are structurally integrated into language as a proportional background. Interestingly, Altmann et al. [41] recently showed that “word niche” can be an essential factor in modeling word use dynamics. New niche words, though they are marginal increases to a language’s lexicon, are themselves anything but “marginal” - they are core words within a subset of the language. This is particularly the case in online communities in which individuals strive to distinguish themselves on short timescales by developing stylistic jargon, highlighting how language patterns can be context dependent.

We now return to the relation between Heaps’ law and Zipf’s law. Table I summarizes the b values calculated by means of ordinary least squares regression using $U_c = 0$ to relate $N_u(t)$ to $N_w(t)$. For $U_c = 1$ we find that $b \approx 0.5$ for all languages analyzed, as expected from Heaps law, but for $U_c \gtrsim 8$ the b value significantly deviates from 0.5, and for $U_c \gtrsim 1000$ the b value begins to saturate approaching unity. Considering that $\alpha_+ \approx 2$ implies $\zeta \approx 1$ for all corpora, Figures 3 and 4 shows that we can confirm the relation $b(U_c) \approx 1/\zeta$ only for the more pruned corpora that require relatively large U_c . This hidden feature of the scaling relation highlights the underlying structure of language, which forms a dependency network between the common words of the kernel lexicon and their more esoteric counterparts in the unlimited lexicon. Moreover, the function $\partial N_w / \partial N_u \sim (N_u)^{b-1}$ is a monotonically decreasing function for $b < 1$, demonstrating the *decreasing marginal need* for additional words as a corpora grows. In other words, since we get more and more “mileage” out of new words in an already large language, additional words are needed less and less.

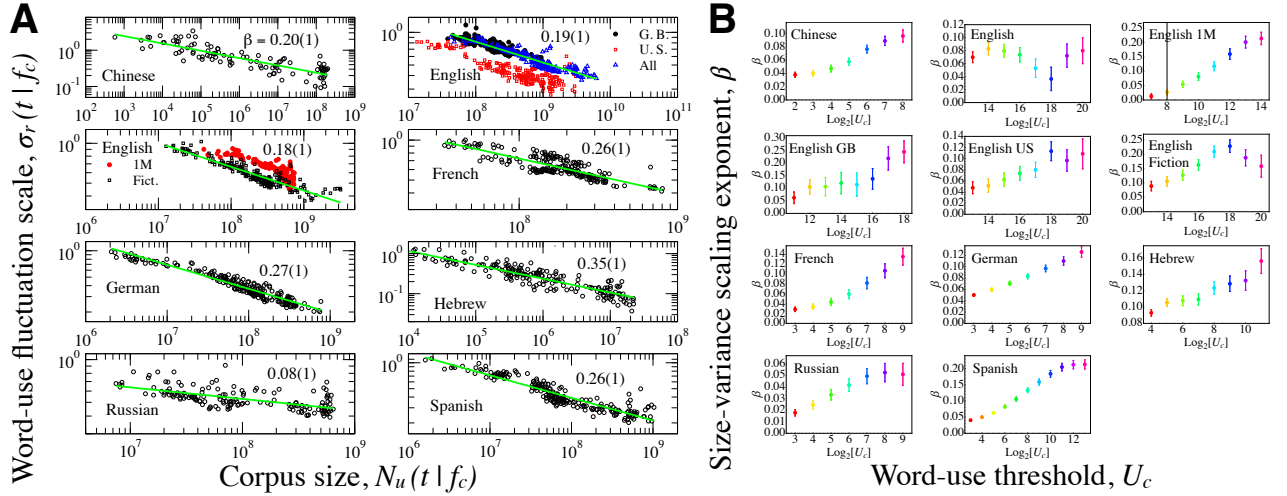


FIG. 7: Growth fluctuation of word use scale with the size of the corpora. (A) Depicted is the quantitative relation in Eq.(8) between $\sigma_r(t|f_c)$ and the corpus size $N_u(t|f_c)$. We calculate $\sigma_r(t|f_c)$ using the relatively common words that meet the criterion that their average word use $\langle f_i \rangle$ over the entire word history (using words with lifetime $T_i \geq 10$ years) is larger than a threshold $f_c \equiv 10/\text{Min}[N_u(t)]$ (see Table I). We show the language-dependent scaling value $\beta \approx 0.08 - 0.35$ in each panel. For each language we show the value of the ordinary least squares best-fit β value with the standard error in parentheses. (B) Summary of $\beta(U_c)$ exponents calculated using a use-threshold U_c , instead of a frequency threshold f_c as used in (A). Error bars indicate the standard error in the OLS regression. We perform this additional analysis in order to provide alternative insight into the role of extremely rare words. For increasing U_c the $\beta(U_c)$ value for each corpora increases from $\beta \approx 0.05$ to $\beta < 0.25$. This language pruning method quantifies the role of new rare words (also including OCR errors, spelling and other orthographic variants), which are the significant components of language volatility.

Corpora size and word-use fluctuations. Lastly, it is instructive to examine how vocabulary size N_w and the overall size of the corpora N_u affect fluctuations in word use. Figure 5 shows how $N_w(t)$ and $N_u(t)$ vary over time over the past two centuries. Note that, apart from the periods during the two World Wars, the number of words printed, which we will refer to as the “literary productivity”, has been increasing over time. The number of distinct words (vocabulary size) has also increased reflecting basic social and technological advancement [8].

To investigate the role of fluctuations, we focus on the logarithmic growth rate, commonly used in finance and economics

$$r_i(t) \equiv \ln f_i(t + \Delta t) - \ln f_i(t) = \ln \left[\frac{f_i(t + \Delta t)}{f_i(t)} \right], \quad (7)$$

to measure the relative growth of word use over 1-year periods, $\Delta t \equiv 1$ year. Recent quantitative analysis on the distribution $P(r)$ of word use growth rates $r_i(t)$ indicates that annual fluctuations in word use deviates significantly from the predictions of null models for language evolution [9].

We define an aggregate fluctuation scale, $\sigma_r(t|f_c)$, using a frequency cutoff $f_c \propto 1/\text{Min}[N_u(t)]$ to eliminate infrequently used words. The quantity $\text{Min}[N_u(t)]$ is the minimum corpora size over the period of analysis, and so $1/\text{Min}[N_u(t)]$ is an upper bound for the minimum observed frequency for words in the corpora. Figure 6 shows $\sigma_r(t|f_c)$, the standard deviation of $r_i(t)$ calculated across all words that satisfy the condition $\langle f_i \rangle \geq f_c$ for words with lifetime $T_i \geq 10$ years, using $f_c = 1/\text{Min}[N_u(t)]$. Visual inspection

suggests a general decrease in $\sigma_r(t|f_c)$ over time, marked by sudden increases during times of political conflict. Hence, the persistent increase in the volume of written language is correlated with a persistent downward trend what could be thought of as the “system temperature” $\sigma_r(t|f_c)$: as a language grows and matures it also “cools off”.

Since this cooling pattern could arise as a simple artifact of an independent identically distributed (i.i.d) sampling from an increasingly large dataset, we test the scaling of $\sigma_r(t|f_c)$ with corpora size. Figure 7(A) shows that for large $N_u(t)$, each language is characterized by a scaling relation

$$\sigma_r(t|f_c) \sim N_u(t|f_c)^{-\beta}, \quad (8)$$

with language-dependent scaling exponent $\beta \approx 0.08 - 0.35$. We use $f_c = 10/\text{Min}[N_u(t)]$, which defines the frequency threshold for the inclusion of a given word in our analysis. There are two candidate null models which give insight into the limiting behavior of β . The Gibrat proportional growth model predicts $\beta = 0$ and the Yule-Simon urn model predicts $\beta = 1/2$ [42]. We observe $\beta < 1/2$, which indicates that the fluctuation scale decreases more slowly with increasing corpora size than would be expected from the Yule-Simon urn model prediction, deducible via the “delta method” for determining the approximate scaling of a distribution and its standard deviation σ [43].

To further compare the roles of the kernel lexicon versus the unlimited lexicon, we apply our pruning method to quantify the dependence of the scaling exponent β on the fluctuations arising from rare words. We omit words from our calculation of $\sigma_r(t|U_c)$ if their use $u_i(t)$ in year t falls below the

word-use threshold U_c . Fig. 7(B) shows that $\beta(U_c)$ increases from values close to 0 to values less than $1/2$ as U_c increases exponentially. An increasing $\beta(U_c)$ confirms our conjecture that rare words are largely responsible for the fluctuations in a language. However, because of the dependency structure between words, there are residual fluctuation spillovers into the kernel lexicon likely accounting for the fact that $\beta < 1/2$ even when the fluctuations from the unlimited lexicon are removed.

A size-variance relation showing that larger entities have smaller characteristic fluctuations was also demonstrated at the scale of individual words using the same *Google n-gram* dataset [9]. Moreover, this size-variance relation is strikingly analogous to the decreasing growth rate volatility observed as complex economic entities (i.e. firms or countries) increase in size [42, 44–48], which strengthens the analogy of language as a complex ecosystem of words governed by competitive forces.

Further possible explanations for $\beta < 1/2$ is that language growth is counteracted by the influx of new words which tend to have growth-spurts around 30-50 years following their birth in the written corpora [9]. Moreover, the fluctuation scale $\sigma_r(t|f_c)$ is positively influenced by adverse conditions such as wars and revolutions, since a decrease in $N_u(t)$ may decrease the competitive advantage that old words have over new words, allowing new words to break through. The globalization effect, manifesting from increased human mobility during periods of conflict, is also responsible for the emergence of new words within a language.

Discussion

A coevolutionary description of language and culture requires many factors and much consideration [49, 50]. While scientific and technological advances are largely responsible for written language growth as well as the birth of many new words [9], socio-political factors also play a strong role. For example, the sexual revolution of the 1960s triggered the sudden emergence of the words “girlfriend” and “boyfriend” in the English corpora [1], illustrating the evolving culture of romantic courting. Such technological and socio-political perturbations require case-by-case analysis for any deeper understanding, as demonstrated comprehensively by Michel et al. [8].

Here we analyzed the macroscopic properties of written language using the *Google Books* database [1]. We find that the word frequency distribution $P(f)$ is characterized by two scaling regimes. While frequently used words that constitute the kernel lexicon follow the Zipf law, the distribution has a less-steep scaling regime quantifying the rarer words constituting the *unlimited lexicon*. Our result is robust across languages as well as across other data subsets, thus extending the validity of the seminal observation by Ferrer i Cancho and Solé [14], who first reported it for a large body of English text. The kink in the slope preceding the entry into the unlimited lexicon is a likely consequence of the limits of human mental ability that force the individual to optimize the usage

of frequently used words and forget specialized words that are seldom used. This hypothesis agrees with the “principle of least effort” that minimizes communication noise between speakers (writers) and listeners (readers), which in turn may lead to the emergence of the Zipf law [16].

Using an extremely large written corpora that documents the profound expansion of language over centuries, we analyzed the dependence of vocabulary growth on corpus growth and validate the Heaps law scaling relation given by Eq. 5. Furthermore we systematically prune the corpora data using a word occurrence threshold U_c , and comparing the resulting $b(U_c)$ value to the $\zeta \approx 1$ value, which is stable since it is derived from the “kernel” lexicon. We conditionally confirm the theoretical prediction $\zeta \approx 1/b$ [13, 34–38], which we validate only in the case that the extremely rare “unlimited” lexicon words are not included in the data sample (see Figs. 3 and 4).

The economies of scale ($b < 1$) indicate that there is an *increasing marginal return* for new words, or alternatively, a *decreasing marginal need* for new words, as evidenced by allometric scaling. This can intuitively be understood in terms of the increasing complexities and combinations of words that become available as more words are added to a language, lessening the need for lexical expansion. However, a relationship between new words and existing words is retained. Every introduction of a word, from an informal setting (e.g. an expository text) to a formal setting (e.g. a dictionary) is yet another chance for the more common describing words to play out their respective frequencies, underscoring the hierarchy of words. This can be demonstrated quite instructively from Eq. (6) which implies that for $b = 1/2$ that $\frac{\partial N_u}{\partial N_w} \propto N_w$, meaning that it requires a quantity proportional to the vocabulary size N_w to introduce a new word, or alternatively, that a quantity proportional to N_w necessarily results from the addition.

Though new words are needed less and less, the expansion of language continues, doing so with marked characteristics. Taking the growth rate fluctuations of word use to be a kind of temperature, we note that like an ideal gas, most languages “cool” when they expand. The fact that the relationship between the temperature and corpus volume is a power law, one may, loosely speaking, liken language growth to the expansion of a gas or the growth of a company [42, 44–48]. In contrast to the static laws of Zipf and Heaps, we note that this finding is of a dynamical nature.

Other aspects of language growth may also be understood in terms of expansion of a gas. Since larger literary productivity imposes a downward trend on growth rate fluctuations — which also implies that the ranking of the top words and phases becomes more stable [51] — productivity itself can be thought of as a kind of inverse pressure in that highly productive years are observed to “cool” a language off. Also, it is during the “high-pressure” low productivity years that new words tend to emerge more frequently.

Interestingly, the appearance of new words is more like gas condensation, tending to cancel the cooling brought on by language expansion. These two effects, corpus expansion and new word “condensation,” therefore act against each other. Across all corpora we calculate a size-variance scaling ex-

ponent $0 < \beta < 1/2$, bounded by the prediction of $\beta = 0$ (Gibrat growth model) and $\beta = 1/2$ (Yule-Simon growth model) [42].

In the context of allometric relations, Bettencourt et al. [27] note that the scaling relations describing the dynamics of cities show an *increase* in the characteristic pace of life as the system size grows, whereas those found in biological systems show *decrease* in characteristic rates as the system size grows.

Since the languages we analyzed tend to “cool” as they expand, there may be deep-rooted parallels with biological systems based on principles of efficiency [16]. Languages, like biological systems demonstrate economies of scale ($b < 1$) manifesting from a complex dependency structure that mimics a hierarchical “circulatory system” required by the organization of language [39, 52–56] and the limits of the efficiency of the speakers/writers who exchange the words [19, 41, 57].

-
- [1] Google Books Ngram Viewer, <http://books.google.com/ngrams>.
 - [2] Evans, J. A. and Foster, J. G. Metaknowledge. *Science* **331**, 721–725 (2011).
 - [3] Ball, P. *Why Society is a Complex Matter*. (Springer-Verlag, Berlin, 2012).
 - [4] Helbing, D., Battiéti, S. How to Create an Innovation Accelerator. *Eur. Phys. J. Special Topics* **195**, 101–136 (2011).
 - [5] Lazer, D., et al. Computational social science. *Science* **323**, 721–723 (2009).
 - [6] Barabási, A. L. The network takeover. *Nature Physics* **8**, 14–16 (2012).
 - [7] Vespignani, A. Modeling dynamical processes in complex socio-technical systems. *Nature Physics* **8**, 32–39 (2012).
 - [8] Michel, J.-B., et al. Quantitative analysis of culture using millions of digitized books. *Science* **331**, 176–182 (2011).
 - [9] Petersen, A. M., Tenenbaum, J., Havlin, S., and Stanley, H. E. Statistical laws governing fluctuations in word use from word birth to word death. *Scientific Reports* **2**, 313 (2012).
 - [10] Gao, J., Hu, J., Mao, X., and Perc, M. Culturomics meets random fractal theory: Insights into long-range correlations of social and natural phenomena over the past two centuries. *J. R. Soc. Interface* **9**, 1956–1964 (2012).
 - [11] Zipf, G. K. *Human Behavior and the Principle of Least-Effort: An Introduction to Human Ecology*. Addison-Wesley, Cambridge, MA, (1949).
 - [12] Tsonis, A. A., Schultz, C., and Tsonis, P. A. Zipf’s law and the structure and evolution of languages. *Complexity* **3**, 12–13 (1997).
 - [13] Serrano, M. Á., Flammini, A., and Menczer, F. Modeling statistical properties of written text. *PLoS ONE* **4**, e5372 (2009).
 - [14] Ferrer i Cancho, R. and Solé, R. V. Two regimes in the frequency of words and the origin of complex lexicons: Zipf’s law revisited. *Journal of Quantitative Linguistics* **8**, 165–173 (2001).
 - [15] Ferrer i Cancho, R. The variation of Zipf’s law in human language. *Eur. Phys. J. B* **44**, 249–257 (2005).
 - [16] Ferrer i Cancho, R. and Solé, R. V. Least effort and the origins of scaling in human language. *Proc. Natl. Acad. Sci. USA* **100**, 788–791 (2003).
 - [17] Baek, S. K., Bernhardsson, S., and Minnhagen, P. Zipf’s law unzipped. *New J. Phys.* **13**, 043004 (2011).
 - [18] Heaps, H. S. *Information Retrieval: Computational and Theoretical Aspects*. (Academic Press, New York, 1978).
 - [19] Bernhardsson, S., Correa da Rocha, L. E., and Minnhagen, P. The meta book and size-dependent properties of written language. *New J. Phys.* **11**, 123015 (2009).
 - [20] Bernhardsson, S., Correa da Rocha, L. E., and Minnhagen, P. Size-dependent word frequencies and translational invariance of books. *Physica A* **389**, 330–341 (2010).
 - [21] Kleiber, M. Body size and metabolism. *Hilgardia* **6**, 315–351 (1932).
 - [22] West, G. B. Allometric scaling of metabolic rate from molecules and mitochondria to cells and mammals. *Proc. Natl. Acad. Sci. USA* **98**, 2473–2478 (2002).
 - [23] Makse, H. A., Havlin, S., and Stanley, H. E. Modelling urban growth patterns. *Nature* **377**, 608–612 (1995).
 - [24] Makse, H. A., Jr., J. S. A., Batty, M., Havlin, S., and Stanley, H. E. Modeling urban growth patterns with correlated percolation. *Phys. Rev. E* **58**, 7054–7062 (1998).
 - [25] Rozenfeld, H. D., Rybski, D., Andrade Jr., J. S., Batty, M., Stanley, H. E., and Makse, H. A. Laws of population growth. *Proc. Natl. Acad. Sci. USA* **48**, 18702–18707 (2008).
 - [26] Gabaix, X. Zipf’s law for cities: An explanation. *Quarterly Journal of Economics* **114**, 739–767 (1999).
 - [27] Bettencourt, L. M. A., Lobo, J., Helbing, D., Kuhnert, C., and West, G. B. Growth, innovation, scaling, and the pace of life in cities. *Proc. Natl. Acad. Sci. USA* **104**, 7301–7306 (2007).
 - [28] Batty, M. The size, scale, and shape of cities. *Science* **319**, 769–771 (2008).
 - [29] Rozenfeld, H. D., Rybski, D., Gabaix, X., and Makse, H. A. The area and population of cities: New insights from a different perspective on cities. *American Economic Review* **101**, 2205–2225 (2011).
 - [30] Newman, M. E. J. Power laws, Pareto distributions and Zipf’s law. *Contemporary Phys.* **46**, 323–351 (2005).
 - [31] Stanley, M. H. R., Buldyrev, S. V., Havlin, S., Mantegna, R., Salinger, M., and Stanley, H. E. Zipf plots and the size distribution of firms. *Econ. Lett.* **49**, 453–457 (1995).
 - [32] Mantegna, R. N., et al. Systematic analysis of coding and non-coding DNA sequences using methods of statistical linguistics. *Phys. Rev. E* **52**, 2939–2950 (1995).
 - [33] Clauset, A., Shalizi, C. R., and Newman, M. E. J. Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703 (2009).
 - [34] Mandelbrot, B. On the theory of word frequencies and on related Markovian models of discourse, in: R. Jakobson, Structure of Language and its Mathematical Aspects. *Proceedings of Symposia in Applied Mathematics* Vol. **XII**, 190–219 (1961).
 - [35] Karlin, S. Central limit theorems for certain infinite urn schemes. *Journal of Mathematics and Mechanics* **17**, 373–401 (1967).
 - [36] Gneden, A., Hansen, B., Pitman, J. Notes on the occupancy problem with infinitely many boxes: general asymptotics and power laws. *Probability Surveys* **4**, 146–171 (2007).
 - [37] van Leijenhorst, D. C., van der Weide, Th. P. A formal derivation of Heaps’ Law. *Inform. Sci.* **170**, 263–272 (2005).
 - [38] Lü, L., Zhang, Z.-K., Zhou, T. Zipf’s law leads to Heaps’ law: Analyzing their relation in finite-size systems. *PLoS One* **5**, e14139 (2010).
 - [39] Steyvers, M. and Tenenbaum, J. B. The large-scale structure of semantic networks: Statistical analyses and a model of semantic growth. *Cogn. Sci.* **29**, 41–78 (2005).

- [40] Markosova, M. Network model of human language. *Physica A* **387**, 661–666 (2008).
- [41] Altmann, E. G., Pierrehumbert, J. B., and Motter, A. E. Niche as a determinant of word fate in online groups. *PLoS ONE* **6**, e19009 (2011).
- [42] Riccaboni, M., Pammolli, F., Buldyrev, S. V., Ponta, L., and Stanley, H. E. The size variance relationship of business firm growth rates. *Proc. Natl. Acad. Sci. USA* **105**, 19595–19600 (2008).
- [43] Oehlert, G. W. A Note on the Delta Method. *The American Statistician* **46**, 27–29 (1992).
- [44] Amaral, L. A. N., et al. Scaling Behavior in Economics: I. Empirical Results for Company Growth. *J. Phys. I France* **7**, 621–633 (1997).
- [45] Amaral, L. A. N., et al. Power Law Scaling for a System of Interacting Units with Complex Internal Structure. *Phys. Rev. Lett.* **80**, 1385–1388 (1998).
- [46] Fu, D., Pammolli, F., Buldyrev, S. V., Riccaboni, M., Matia, K., Yamasaki, K., and Stanley, H. E. The growth of business firms: Theoretical framework and empirical evidence. *Proc. Natl. Acad. Sci. USA* **102**, 18801–18806 (2005).
- [47] Podobnik, B., Horvatic, D., Petersen, A. M., Stanley, H. E. Quantitative relations between risk, return, and firm size. *EPL* **85**, 50003 (2009).
- [48] Podobnik, B., Horvatic, D., Petersen, A. M., Njavro, M., Stanley, H. E. Common scaling behavior in finance and macroeconomics. *Eur. Phys. J. B* **76**, 487–490 (2010). *EPL* **85**, 50003 (2009).
- [49] Mufwene, S. *The Ecology of Language Evolution*. (Cambridge Univ. Press, Cambridge, UK, 2001).
- [50] Mufwene, S. *Language Evolution: Contact, Competition and Change*. (Continuum International Publishing Group, New York, NY, 2008).
- [51] Perc, M. Evolution of the most common English words and phrases over the centuries. *J. R. Soc. Interface* **9**, 3323–3328 (2012).
- [52] Sigman, M. and Cecchi, G. A. Global organization of the word-net lexicon. *Proc. Natl. Acad. Sci. USA* **99**, 1742–1747 (2002).
- [53] Alvarez-Lacalle, E., Dorow, B., Eckmann, J.-P., and Moses, E. Hierarchical structures induce long-range dynamical correlations in written texts. *Proc. Natl. Acad. Sci. USA* **103**, 7956–7961 (2006).
- [54] Altmann, E. A., Cristadoro, G., and Esposti, M. D. On the origin of long-range correlations in texts. *Proc. Natl. Acad. Sci. USA* **109**, 11582–11587 (2012).
- [55] Montemurro, M. A. and Pury, P. A. Long-range fractal correlations in literary corpora. *Fractals* **10**, 451–461 (2002).
- [56] Corral, A., Ferrer i Cancho, R., and Díaz-Guilera, A. Universal complex structures in written language. *arXiv:0901.2924* (2009).
- [57] Altmann, E. G., Pierrehumbert, J. B., and Motter, A. E. Beyond word frequency: bursts, lulls, and scaling in the temporal distributions of words. *PLoS ONE* **4**, e7678 (2009).

Acknowledgments

AMP acknowledges support from the IMT Lucca Foundation. JT, SH and HES acknowledge support from the DTRA, ONR, the European EPIWORK and LINC projects, and the Israel Science Foundation. MP acknowledges support from the Slovenian Research Agency.

Author Contributions

A. M. P., J. T., S. H., H. E. S., & MP designed research, performed research, wrote, reviewed and approved the manuscript. A. M. P. performed the numerical and statistical analysis of the data.